

C Level 1 — The Common Core: Clinical AI Foundations

A step-by-step curriculum blueprint for the 12-hour foundational training.

Curriculum Vital Statistics



Duration
12 contact hours
(2 days of 6, or 3 of 4)



Audience
Mixed Cadre
(16–24 learners)



Instructors
2 Level-4 Certified Facilitators



Environment
1 Controlled Sandbox
(Interaction logging active)

FOUNDATIONAL DESIGN DECISIONS

Establishing the non-negotiable rules of the clinical training environment.



Mixed Cadre, Deliberately.

Surgeon and ward nurse sit in the same room. AI failure modes live in the space between cadres. Nurses must practice escalating concerns; consultants must practice expecting challenges. Segregated cohorts produce a smoother day but a worse course.



Controlled Sandbox Only.

Personal devices are strictly prohibited.

Provides a safe environment to log interactions, utilize seeded-error configurations for testing, and enforce a zero-tolerance policy on real patient data entry.



Fixed Sequencing.


The order of units is locked.

Learners must understand system realities and failure modes before practicing discernment, and they must know the law before touching real clinical workflows.

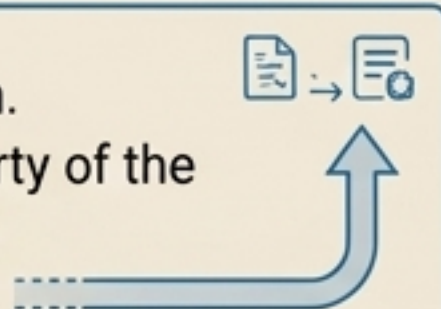


FACILITATOR NOTE: Emphasize the non-negotiable nature of these foundational decisions. Resistance from participants regarding mixed cohorts or fixed sequencing must be addressed immediately, linking back to the core objective of simulating real-world clinical AI integration and its inherent challenges.

⚠️ CONJUNCTIVE FAILURE TRAP: Permitting personal device use for 'convenience' or allowing re-sequencing due to participant preference critically undermines the controlled environment and invalidates performance data.

THE 12-HOUR CURRICULUM ARCHITECTURE



Unit 1.1 — Demystifying the Systems

	System Type	Core Function	Trust Level
1	Language Model	Next-token prediction. Plausibility is a property of the text, not of the world. Fluency ≠ Accuracy. 	Requires independent verification. 
2	Diagnostic Classifier	Probability estimation from trained models using patient parameters. 	Evaluated against specific clinical thresholds. 
3	Rule-based Alert	Triggered deterministic thresholds (e.g., flagging gentamicin prescriptions when creatinine exceeds a limit). 	Deterministic. 



FACILITATOR NOTE: The Fabricated Citation Live Demo

Instructors ask the sandbox for evidence on a local Kenyan clinical question [] and request references. []
FABRICATED CITATION.
FABRICATED

Live lookup proves that fluent, correctly formatted citations are often entirely fabricated, establishing immediately that fluency does not equal truth.

Unit 1.2 – The Three Modalities of AI Use

1	Automation	2	Augmentation	3	Agency
DEFINITION	---	DEFINITION	---	DEFINITION	---
	Unreviewed output entering a record.		Human actively working with AI.		AI executing tasks alone.
EXAMPLE	---	EXAMPLE	---	EXAMPLE	---
	A midwife asking a system to translate a pre-eclampsia leaflet into Kiswahili.		A physician reasoning through a differential diagnosis with the system.		A triage assistant configured to pre-sort an outpatient queue.
RISK	---	RISK	---	RISK	---
	Unchecked translation quality errors.		Cognitive anchoring bias.		Harm at scale with no human in the loop when an error occurs.


SYNTHESIS INSIGHT

THE CATASTROPHIC RISK: Most critical failures occur when a clinician operates in Agency mode while believing they are in Augmentation mode.

Unit 1.2 — The "Two Loops" Framework

Mapping AI competencies directly onto familiar clinical workflows

The Legal Frame
(Consent & Audit)




The Legal Frame
(Consent & Audit)

SYNTHESIS

AI competencies are not foreign technical skills; they are modern extensions of consent, audit, history, and examination.

Unit 1.3 — Delegation & The Reversibility Test



Core Principle

Reversibility and severity, not convenience or workload, dictate delegability. If the output is wrong and nobody notices, what is the cost?

The Non-Delegable List


- Obtaining consent
- Breaking significant news
- Decision to operate
- Writing the prescription
- Signing the record

FACILITATOR INSIGHT

Cross-cadre pairings are vital. Nurses typically write absolute lists, while doctors write conditional lists. Debating these boundaries teaches more than the lists themselves.

Unit 1.4 — Description as Clinical Handover

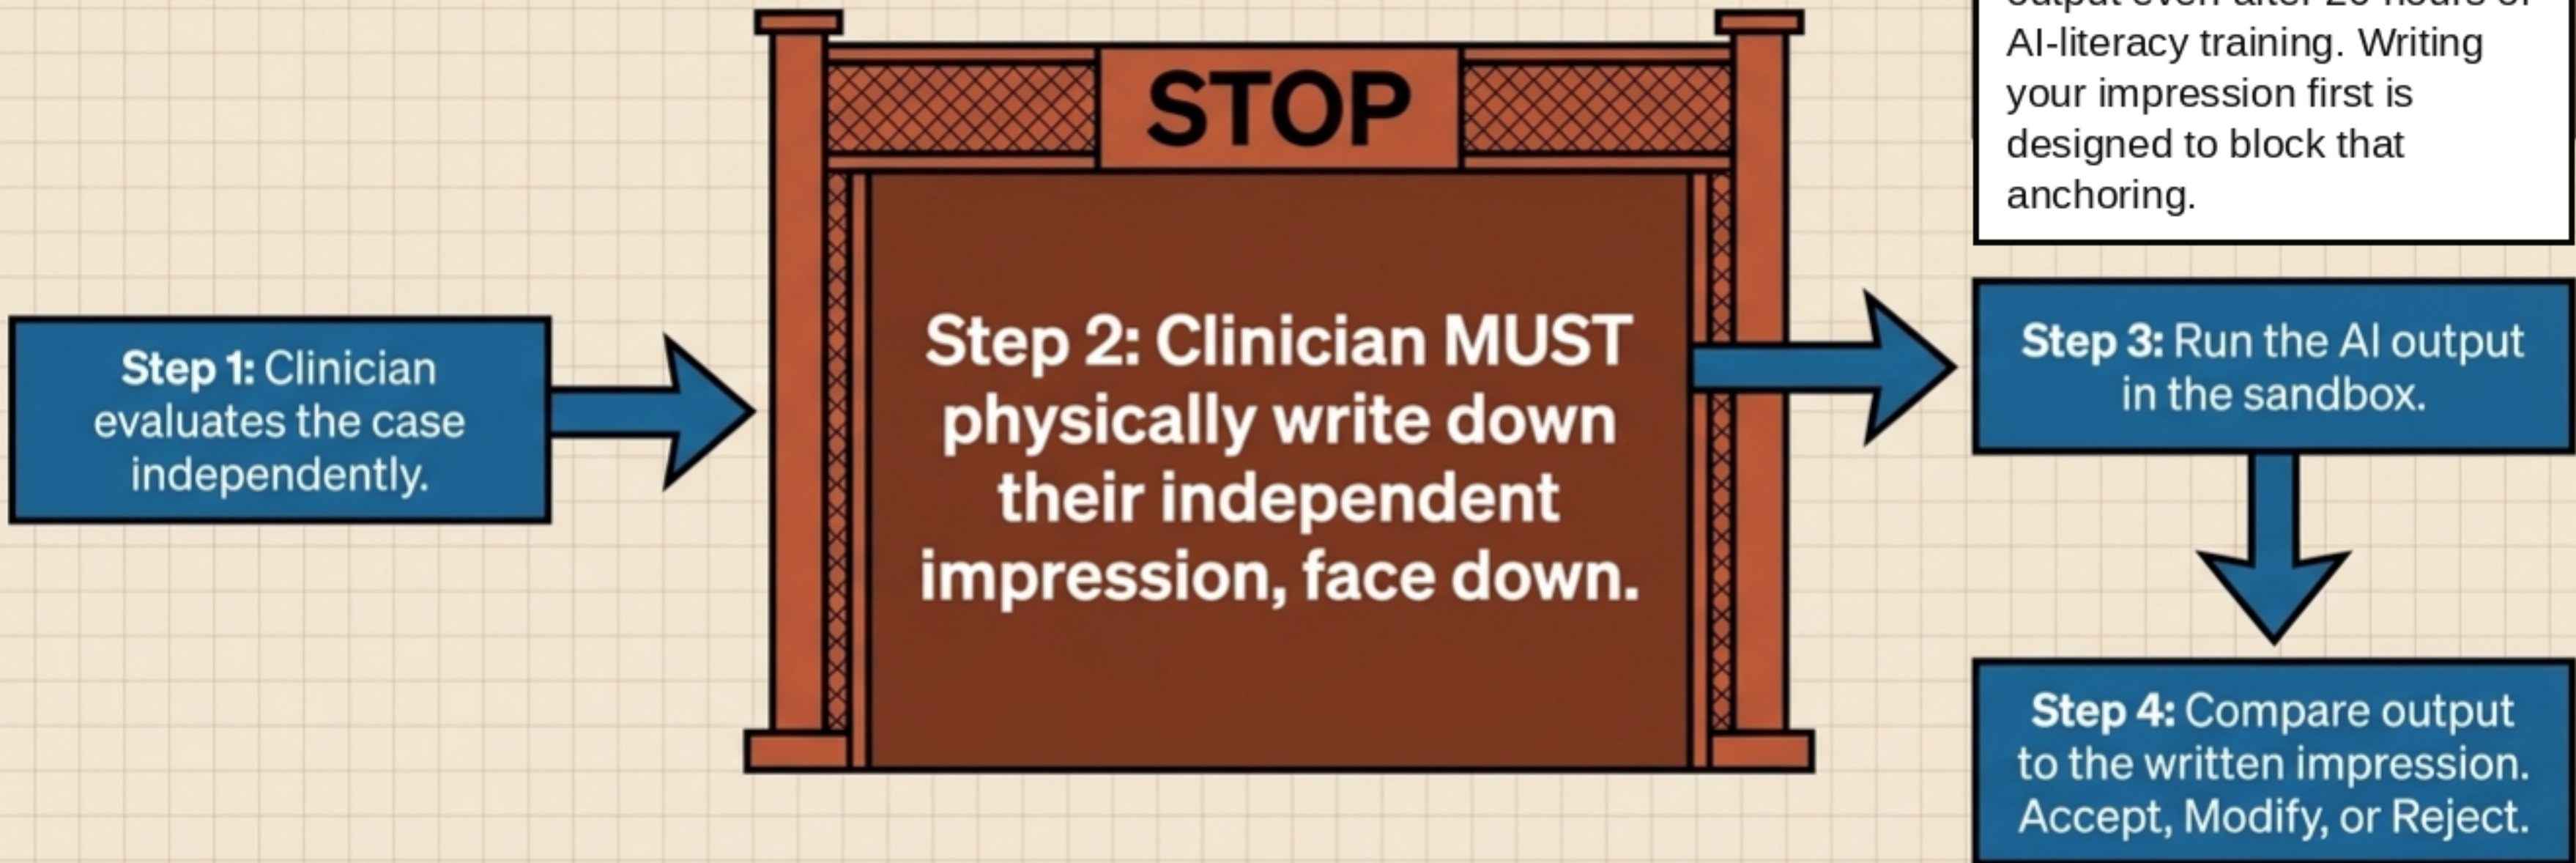
Perfect AI Prompt



Conjunctive Failure Trap

ZERO IDENTIFIABLE DATA. The absolute rule. Entering identifiable patient data into a system you do not control results in an automatic **course failure**.

Unit 1.5 — Discernment & The Independent-Impression Rule




The Evidence

In 2025 trials, clinicians deferred to erroneous AI output even after 20 hours of AI-literacy training. Writing your impression first is designed to block that anchoring.

Core Principle

Pre-recording the human impression makes the clinician's reasoning a fixed point, preventing quiet cognitive anchoring.

Unit 1.5 — Seeded Errors & Calibrated Trust



THE SANDBOX DRILL

To build reflex detection, learners face cases with a 1-in-3 seeded error rate.



Example Scenario:

An AI differential for a 29-year-old with weight loss and night sweats that fluently omits Tuberculosis.

Local clinical knowledge is the only safeguard.

Unit 1.6 — Law, Consent, and Documentation

The Law

Governed by the
**Data Protection
Act (2019) &
Digital Health Act (2023).**




**Identifiable data
stays inside systems
you or your employer
control.**

The Patient

Disclosure role-play
demands
honest framing.



*A system was used to
help think it through,
I reviewed it, and the
decision is mine.*

(Neither a simple 'yes' nor 'no').

The Record

Standardized
documentation is
strictly required.



*Differential generated
with AI decision
support; reviewed and
modified by me; final
plan is my own.*

**THE BOTTOM LINE: THE SIGNATURE IS YOURS.
LIABILITY CANNOT BE DELEGATED TO AN ALGORITHM.**


Unit 1.7 — Synthesis: The Personal AI Policy

The physical artifact every learner creates and takes back to their practice.



 **PERSONAL AI POLICY**

- 1.**
Approved uses in my specific clinical practice.
- 2.**
My strict non-delegable list (from Unit 1.3).
- 3.**
Data entry constraints (from Unit 1.4).
- 4.**
Mechanisms for checking output before acting (from Unit 1.5).
- 5.**
Standardized documentation phrasing (from Unit 1.6).
- 6.**
Escalation procedures for AI-contributed harm or near-misses.
- 7.**
Mandatory review date.

X  
Clinician Signature & Date

Assessment Architecture

Knowledge Assessment



40-item invigilated MCQ.



Standard set by modified Angoff panel
(No arbitrary 50% pass rate).



40% of total weight is dedicated entirely
to Discernment (D3).



Simulated Encounter



10-minute observed standardized patient
sandbox encounter.



Conjunctive Failure Trap

AUTOMATIC FAIL: Identifiable Data Breach.
Entering patient info into an uncontrolled system.



Conjunctive Failure Trap

AUTOMATIC FAIL: Failing to detect and
correct the seeded clinical error at threshold.

Final Outcome



Certificate in Clinical AI Foundations.

Mandatory prerequisite for Level 2 supervised use in clinical workflows. Recertification required at 2 years.

